

Improved Photometric Redshifts via Enhanced Estimates of System Response, Galaxy Templates, and Magnitude Priors

Samuel J. Schmidt^{*} and Paul Thorman[†]

Dept. of Physics, University of California, One Shields Ave., Davis, CA 95616

15 November 2012

ABSTRACT

Wide, deep photometric surveys require robust photometric redshift estimates (photo-z's) for studies of large-scale structure. These estimates depend critically on accurate photometry. We describe the improvements to the photometric calibration and the photo-z estimates in the Deep Lens Survey (DLS) from correcting three of the inputs to the photo-z calculation: the system response as a function of wavelength, the spectral energy distribution templates, and template prior probabilities as a function of magnitude. We model the system response with a physical model of the MOSAIC camera's CCD, which corrects a 0.1 magnitude discrepancy in the colours of type M2 and later stars relative to the SDSS z photometry. We provide our estimated z response function for the use of other surveys that used MOSAIC before its recent detector upgrade. The improved throughput curve, template set, and Bayesian prior lead to a 20 per cent reduction in photo-z scatter and a reduction of the bias by a factor of more than two. This paper serves as both a photo-z data release description for DLS and a guide for testing the quality of photometry and resulting photo-z's generally.

Key words: cosmology: observations—methods: data analysis—methods: statistical

1 INTRODUCTION

Due to the time and expense of obtaining spectra of large numbers of faint galaxies, deep photometric surveys have increasingly turned to photometric redshift estimates (photo-z's). For the next generation of all-sky deep surveys, such as LSST (Tyson 2002), Euclid (Laureijs et al. 2010), the Dark Energy Survey (Wester & Dark Energy Survey Collaboration 2005), and Pan-STARRS (Kaiser et al. 2002), photometric redshifts will be essential to the study of large-scale structure, baryon acoustic oscillations, and the evolution of the Universe over time. The improved photo-z's described in this paper, which are based on photometry from the Deep Lens Survey (DLS; Wittman et al. 2002), have enabled the first tomographic measurement of weak lensing magnification (Morrison et al. 2012), made it possible to examine galaxy-mass correlations at large scales (Choi et al. 2012), and been used to establish joint constraints on Ω_M and σ_8 using cosmic shear (Jee et al. 2012).

Fast and robust photo-z software (e.g., BPZ, Hyper-

z; Benítez 2000; Bolzonella et al. 2000) can generate detailed redshift estimates based on measured galaxy fluxes in different passbands. A popular method uses spectroscopic templates based on nearby galaxies (e.g., Coleman et al. 1980, hereafter CWW), redshifted and convolved with user-supplied system response functions to generate theoretical colours for each galaxy spectral type as a function of redshift. These are often combined with a Bayesian prior based on the magnitudes (apparent or absolute) and relative abundances of the various spectral types to determine the relative posterior probability of each type and redshift for a given galaxy's measured fluxes and flux errors.

Good photometry is essential for accurate photo-z's. In any photo-z procedure that fits spectral energy distributions (SEDs), the key assumption is that the input fluxes are sampling light from a single isolated galaxy. Any flux missing in a particular band from improper de-blending, seeing differences, or other unrecognized errors will alter the galaxy's colour, and will very likely translate into an error in the inferred redshift. This was specifically demonstrated for various methods of compensating for seeing in Hildebrandt et al. (2012). In addition, errors in the determination of the system response functions (i.e., filter+mirror+CCD response+atmosphere) will prop-

^{*} Email:sschmidt@physics.ucdavis.edu

[†] Email:thorman@physics.ucdavis.edu

agate into the photo-z calculation through the model flux calculations.

These errors are amplified when “training set” data from the survey itself are used to optimize photo-z performance, which is done in this paper and elsewhere (Budavári et al. 2000; Ilbert et al. 2006). Systematic biases present in the photometry can propagate through the analysis and degrade the photo-z performance. In this paper we will discuss one such example: the incorrect specification of the effective system response for the DLS z -band. After correcting the system response curve, we then discuss modifications of the photo-z template set and Bayesian prior for the DLS.

Details of the DLS photometric calibration can be found in the DLS ubercal paper (Wittman et al. 2012) and the DLS data release paper (Wittman et al. in prep.). The effects of the prior probability and the error distribution have previously been studied for the Deep Lens Survey data (Wittman et al. 2007; Margoniner & Wittman 2008). In this paper, in addition to constructing a new prior, we discuss the diagnosis of photometry problems and a technique to correct the filter response curve (§ 3). In § 4.1 we discuss modifications to SED templates for the photo-z calculation, in § 4.2 we construct a prior based on multiple datasets, and in § 5 we describe the resulting DLS photo-z’s. We conclude and discuss future work in § 6.

2 SUMMARY OF OBSERVATIONS

2.1 Imaging

Our analysis is focused on the Deep Lens Survey, a deep $BVRz$ imaging survey carried out with the MOSAIC and MOSAIC-II cameras at KPNO and CTIO between 1999 and 2006. The primary imaging covered 20 deg^2 in five $2^\circ \times 2^\circ$ fields spaced around the sky. The fields were chosen for their lack of bright stars and galaxies, their appropriate position for split-night observations, and the presence of pre-existing deep spectroscopic surveys, when possible. The R -band images were taken during periods of good seeing ($\text{FWHM} < 0.9''$) with a total integration time of 24 ks, while the B , V , and z bands were integrated for 18 ks apiece. Average 5σ detection limits for point sources in B , V , R , and z are 26.5, 26.5, 27, and 24 magnitudes, respectively (all magnitudes on the Vega system), with some minor variations in depth due to variable seeing. There are also shallower areas around the edges of the dithered pointings (“sub-fields”), which received fewer observations than the sub-field centres. Magnitudes and colours were measured using Colorpro (Coe et al. 2006), which corrects for differences in PSF between filters as well as differences in seeing between stacks of images taken at different times. These magnitudes were used to generate improved flat field maps for the MOSAIC instruments (Wittman et al. 2012) using the ubercal method of Padmanabhan et al. (2008). The complete reduction pipeline and cataloguing procedure will be explained in a companion paper to the DLS public data release (Wittman et al. in prep.).

2.2 Spectroscopy

In addition to measured fluxes, the template adjustment method requires that the redshifts of a fiducial sample be known. The DLS Field 5 (F5) imaging was used to select targets for a subset of the PRImus Multi-object Survey (PRIMUS, Coil et al. 2011). PRIMUS is a low-resolution objective prism slit mask spectroscopic survey using the IMACS instrument on the 6.5-metre Magellan telescope. The instrument has spectroscopic resolution $R \sim 40$, resulting in redshifts precise to $\sigma_z = 0.005/(1+z)$. For F5, PRIMUS is complete to $R = 22.8$ and observed a random sample consisting of 30 per cent of objects with $24.0 > R > 22.8$. We begin with 10,695 objects in F5, and after conservative masking and a cut on quality flag for reliable redshifts ($ZCONF > 2$) we are left with a sample of 9107 galaxies, 703 of which are at $R > 23.3$.

In DLS Field 2 (F2), we made use of targeted redshift observations from the Smithsonian Hectospec Lensing Survey (SHELS, Geller et al. 2005, 2010), which used Hectospec, a moderate-resolution fiber-fed spectrograph on the 6.5-metre MMT, to measure redshifts for galaxies with $R < 20.3$ selected from the DLS R-band imaging. Based on repeated galaxy observations, they estimate their redshift error (stated as $cz(1+z)^{-1}$) to be 27 km s^{-1} for emission line objects and 37 km s^{-1} for absorption line objects, negligible for our purposes.

3 CORRECTING THE THROUGHPUT CURVE

3.1 Identifying problems with photometry

The DLS was originally calibrated, in part, by comparing the colours of pointlike sources measured in the science images to the predicted colours of stars generated by integrating library stellar spectra with the system throughput as a function of wavelength. This measured “stellar locus” was used to correct colour measurements across different fields and subfields of the DLS.

In order to compare the DLS photometry to an independent standard, we make use of the eighth release of the Sloan Digital Sky Survey photometry (SDSS; Aihara et al. 2011); however, the two surveys use different filters and CCDs, so the magnitudes of stars in one survey must be converted to equivalent magnitudes in the other before a proper comparison can be made. The magnitude difference depends both on the response of each system and on the SED of the target, so we focus on the stars of our own Galaxy as targets of approximately known SED. The SDSS colours of various spectral types of star have been calculated and measured directly (Covey et al. 2007), and the $(r-i)$ colour was shown to be a monotonic function of spectral type.

Using $(r-i)$ as a proxy for spectral type, we plot the difference between our DLS z magnitudes and the SDSS z magnitudes for well-detected stars ($\text{S/N} > 10$ in SDSS r , i , and z) as a function of $(r-i)$ in Figure 1. The upper and lower panels in Figure 1 are star samples from fields F2 and F5, which were observed from KPNO and CTIO respectively, so the similarity of the colours indicates that the two MOSAIC cameras had similar response functions at the time of the survey. After spatially matching SDSS objects with high S/N to objects in the DLS catalogue in the

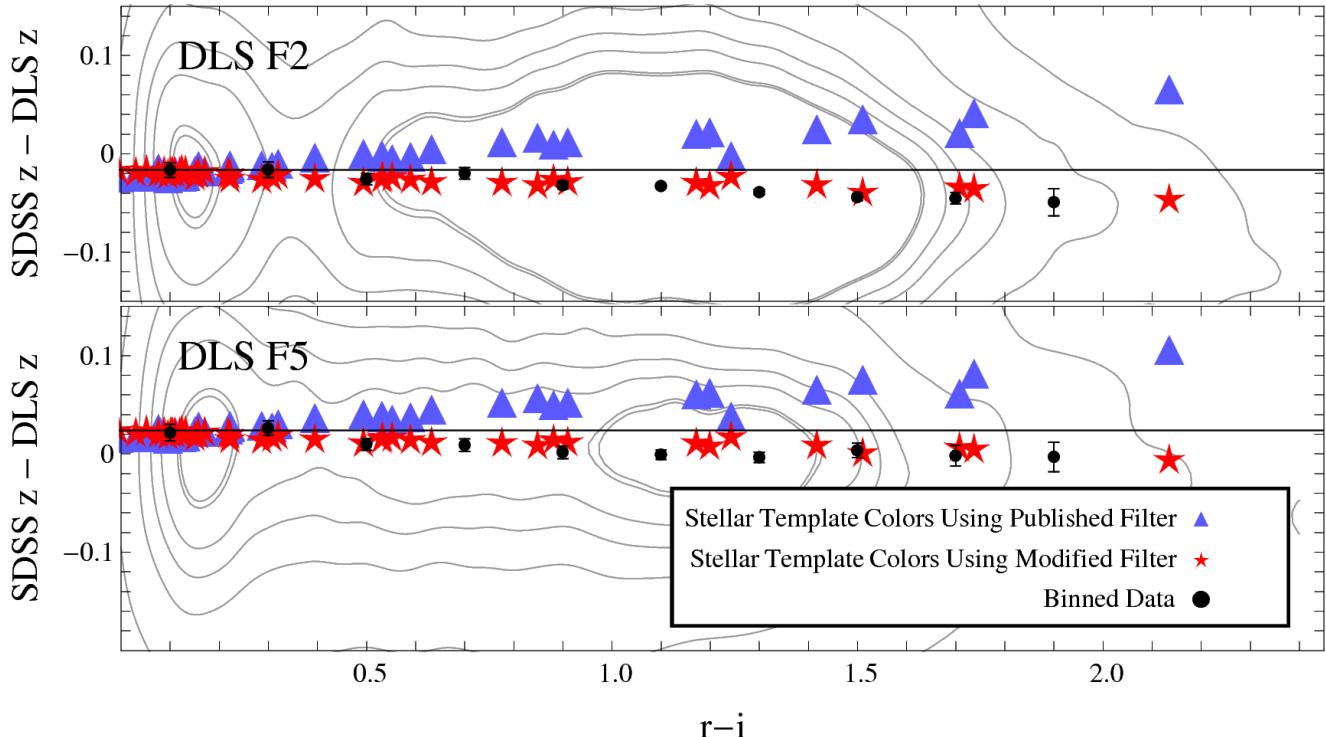


Figure 1. SDSS z -DLS z as a function of SDSS $r - i$ colour. Contours show a kernel density estimator for the colours of the stars, with contours at 0.99, 0.95, 0.68, 0.5, 0.32, 0.05, and 0.01 of the total probability. Mean stellar colours in bins are plotted as black points with error bars. Blue triangles mark the theoretical colours of stars using the published version of the system response function; red stars show the colours with our modified response. F2 and F5 were imaged from KPNO and CTIO respectively, with different MOSAIC cameras, but the modified response function better predicts our measured stellar colours in both cases.

area of overlap (which includes all 4 deg^2 of F2 and approximately 2.63 deg^2 of F5), we select stars by using the DLS catalogue parameter $dlsqcprob$, which is a χ^2 comparison to the measured PSF on a given DLS stack, converted into a likelihood. Using the published system response from the NOAO website¹ and the Pickles (1997) and Bruzual et al. (retrieved 10/06/2009) libraries of stellar spectra, we predict magnitudes for red stars which are significantly brighter than our actual measured magnitudes. The effect is more pronounced for redder stars, and while this colour dependence could be modelled using a classic linear colour term for stars, such a solution could not be expected to extend to galaxies, brown dwarfs, or QSOs, which might share the same ($r - i$) colour but have quite dissimilar detailed spectra. It is therefore necessary to adjust the system response function in order to obtain correct colours for all objects simultaneously.

3.2 Modelling the CCD response

Of the four filters used in the DLS, only the z filter shows signs of differing from its tabulated response curve. The z filters used with MOSAIC and MOSAIC-II were long-pass filters whose transmission reaches 50 per cent of maximum near 8340Å. For incident light of wavelength near $1\mu\text{m}$, the

system response is determined almost entirely by the response of the detector CCDs; the filter is essentially transparent at those wavelengths. The CCDs used in both MOSAIC cameras were thinned, back-illuminated silicon CCDs manufactured by SITe (Scientific Imaging Technologies Inc. 1995), the same detectors used in the Hubble Space Telescope's Advanced Camera for Surveys (Clampin et al. 1998). The absorption of silicon as a function of wavelength and temperature has been modelled for a number of different applications, including solar panels.

Absorption of visible light by silicon occurs by both direct and indirect processes. The direct process excites an electron into the conduction band in a single-photon interaction; the indirect process allows a photon from a wider range of energies to excite the electron in combination with a phonon. The necessity of the phonon in the silicon crystal makes the efficiency of the indirect process strongly temperature-dependent. The phonon can either contribute energy to the interaction or carry it away. We use the parametrization from Rajkhanan et al. (1979), who fitted model coefficients, C_i , to reproduce the direct and indirect band-gap absorption coefficients of silicon, α_{direct} and

¹ <http://www.noao.edu/kpno/mosaic/filters/>

$\alpha_{indirect}$, as a function of temperature, T , in the form

$$\begin{aligned}\alpha(T) &= \alpha_{direct} + \alpha_{indirect} \\ &= A_{direct}(\hbar\omega - E_{direct}(T))^2 \\ &\quad + \sum_{i,j} C_i A_j(T) \left[\frac{\{\hbar\omega - E_{gj}(T) + E_{pi}\}^2}{e^{E_{pi}/k_B T} - 1} \right. \\ &\quad \left. + \rho_i \frac{\{\hbar\omega - E_{gj}(T) - E_{pi}\}^2}{1 - e^{-E_{pi}/k_B T}} \right].\end{aligned}\quad (1)$$

The subscript i runs over possible phonon energies E_{pi} , while the j subscript corresponds to different indirect band gaps E_{gj} . The first term inside the large square brackets corresponds to absorption of a phonon of energy E_{pi} , which process is allowed for any incoming photon whose energy lies in the range $\{E_{gj}(T) - E_{pi}\} \leq \hbar\omega \leq \{E_{gj}(T) + E_{pi}\}$, while the second term represents the emission of a phonon to carry away excess energy; this is only allowed for $\hbar\omega \geq \{E_{gj} + E_{pi}\}$. In Equation 1, the terms must be allowed or disallowed manually, as they do not naturally vanish for forbidden energies. As a result, the function has discontinuities in all derivatives. The temperature dependence of the band gap energies was taken from Varshni (1967), and assumes the form

$$E_g(T) = E_g(0) - \left[\frac{\beta T^2}{T + \gamma} \right]. \quad (2)$$

Rajkanan et al. (1979) then determined the C_i in Equation 1 and the optimal band gap energies by a least-squares fit to the NASA absorption data. In order to convert this α into a quantum efficiency, it is merely necessary to consider the fraction of photons absorbed in a CCD of thickness d :

$$I = I_0 e^{-\alpha d} \quad (\text{definition of } \alpha) \quad (3)$$

$$QE(T, \lambda) = 1 - \frac{I}{I_0} = 1 - e^{-\alpha(T, \lambda)d} \quad (4)$$

We then calculate the quantum efficiency of a silicon detector of the design thickness of the SiTe detectors (14–16 μm) at the operating temperature of the MOSAIC cameras, which is determined from the FITS image headers to be 168K. This is combined with the filter transmission and the atmospheric transmission at KPNO to obtain a modified response function, shown in Figure 2. Atmospheric absorption is based on the atmospheric throughput of Hinkle et al. (2003)², and the overall depth of the absorption features has been allowed to vary as a free parameter in order to optimize the stellar colour fit. The modified MOSAIC z response is slightly less red-sensitive than the SDSS z response; this distinguishes it from the published MOSAIC z response, which claimed a greater red sensitivity than SDSS. This difference in relative sensitivity leads to the qualitative difference in the colour trends of stars in Figure 1. The contours in Figure 1 are a kernel density estimate made by summing a Gaussian error distribution (assumed uncorrelated) for each star. Although there are small discrepancies in zero-point ($\simeq 0.02$ mag) between DLS fields with respect to the predicted stellar locus, these are not relevant to the effect of the z -band response, which is colour-dependent. To remove the effect of the zero-point differences, the predicted SDSS z – DLS z

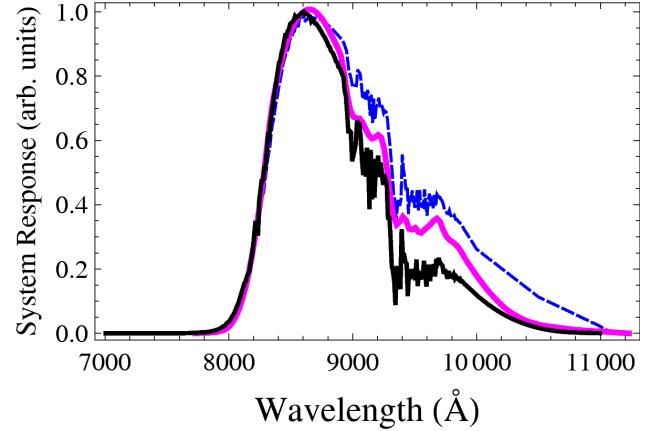


Figure 2. System response functions used in our analysis: published MOSAIC z filter and CCD plus KPNO atmosphere (blue dashed line), published SDSS z response (magenta points), and our modified MOSAIC z response (solid black line). MOSAIC response functions have been averaged over 10 \AA bins for plotting clarity. The modified MOSAIC response at 1 μm has been reduced by 65 per cent from the published value.

colours have been offset by -0.016 in F2 and $+0.024$ in F5 to ensure that zero colour for the library stars matches the mean of the measured star colours for $0 < (r - i) < 0.4$. The errors on the binned stellar colours are derived by first resampling each star from its own error distribution, then bootstrap resampling those stars to determine the standard deviation of the mean; the sampling was carried out 10,000 times on each of 10,000 error distribution resamples. This system response has now been adopted for all of the DLS stellar colour calibrations, as well as the photometric redshift estimation.

This correction to the MOSAIC z -band system response is essential to stellar science, stellar locus calibration, and photo-z's that use this filter. Using the published filter curves leads to disagreement between the blue and red stars of order 0.1 magnitudes in the zero-point of the z filter, which leads to calibration instabilities several times larger than our typical systematic zero-point errors (~ 0.02 mag). For cool star and brown dwarf science, the differences are even more stark: for library L dwarf spectra (the sample from Table 2 of Ryan et al. 2011), the difference between the predicted colours using the published and modified sensitivities is as large as 0.35 magnitudes. The effect of a 0.1 magnitude error on photo-z's is likewise dramatic: as we will show in § 4.1, a 0.1 magnitude miscalibration of the z -band zero-point could lead to a systematic shift of 0.2 in redshift at fixed type; and worse, in practice the types would not be held fixed, leading to substantial misclassification and errors in the template tweaking process. We recommend the adoption of this response curve by anyone interested in using DLS z -band data, or by any other group using archival z -band data obtained with the original MOSAIC cameras.

4 PHOTOMETRIC REDSHIFTS

With a corrected z response function and proper photometry in hand, we now turn to computing the photo-z's. With

² ftp://ftp.noao.edu/catalogs/atmospheric_transmission

only four optical bands to constrain SEDs, we do not expect precise redshift estimates (the survey design goal was ~ 10 per cent uncertainty in photo-z at median redshift). We use version 1.99.3 of the publicly available³ code BPZ (Benítez 2000), a template-based code with Bayesian priors on the apparent magnitude, to perform the redshift estimates. We make adjustments to two components of the procedure in order to optimize photometric redshift performance: the template set used to calculate the initial likelihoods, and the Bayesian apparent magnitude and type prior applied for the final posterior analysis.

One challenge in this procedure (which is also a factor for both the template modification and the prior determination) is that we must perform an initial classification of each galaxy in the training set to a particular galaxy type. We do this by finding the “best-fitted” template to the broad band fluxes at the fixed spectroscopic redshift when compared to those expected from an initial discrete set of SEDs. This is problematic for two reasons: first, because galaxies do not easily separate into discrete types, showing instead a more continuous distribution of properties; and second, because evolution of the galaxy population as a function of redshift may further blur the partitioning of our training sample, and hence our resulting template and prior results. In nearly all current photometric redshift implementations the templates are assumed not to evolve, and model fluxes are determined by simply redshifting the static SED. In this paper we will assume a discretized template set and no evolution in the templates, but we discuss possible modifications in Section 6. For both the template adjustment and prior training, we find the best-fitting type using the built-in ONLY_TYPE setting of the BPZ package. This option constrains the redshift to the correct spectroscopic value and finds the best-fitting template given the four observed DLS fluxes. Note that because we fix the redshift to the correct value, the best-fitting type may be different from the type assigned in the full photo-z fit, which allows the redshift to vary across the full redshift range when determining the type.

4.1 Tweaking the SED Templates

Following Benítez (2000), the heart of template-based photometric redshift codes is a simple χ^2 calculation comparing the observed fluxes of a galaxy with model fluxes found by convolving a model SED with the system throughput with N passbands:

$$\chi^2 = \sum_{i=1}^N \frac{(f_{O,i} - \alpha f_{T,i}(z))^2}{\sigma_i^2} \quad (5)$$

where $f_{O,i}$ is the observed flux in the i th filter, $f_{T,i}(z)$ is the convolved flux of the SED model template at redshift z , σ_i is the flux error, and α is a scaling factor. Note that the use of the α factor to scale all fluxes implicitly assumes neither luminosity nor redshift dependence of galaxy SEDs, only a simple scaling – an assumption that is very likely wrong in practice. We will discuss this further in section 6.

³ <http://www.its.caltech.edu/~coe/BPZ/>

If we assume that our template set roughly spans all possible galaxy types, and we test on a grid of redshifts that covers all possible redshifts, then we can assign a likelihood $\mathcal{L} \propto \exp(-\chi^2/2)$ and normalize the total likelihood to one. We then marginalize over galaxy type (by simply summing the individual type likelihoods) to calculate the one-dimensional redshift probability density function (PDF), $p(z)$. The choice of model SED set is of paramount importance in this calculation. The choice of SED set directly determines the PDF returned by the algorithm. Any differences between the model SEDs chosen to represent the galaxy population and the true population will manifest as a bias in the predicted redshift distribution. We can adjust our template SEDs to better match a set of training data, a process that is often referred to as “tweaking” the templates.

The template correction procedure is straightforward, and similar to the procedure described briefly in Ilbert et al. (2006). We take advantage of the fact that our fixed filter curves sample different wavelength ranges of an SED at different redshifts. In fact, given a set of galaxies with known spectroscopic redshift and classified by SED, we can effectively fix the SED at redshift $z = 0$ and blueshift the filter curves to check which portion of the SED is being sampled. We calculate the effective wavelength, λ_{eff} , of each filter as the weighted mean of the filter transmission curve $T(\lambda)$:

$$\lambda_{eff} = \int \lambda T(\lambda) d\lambda \quad (6)$$

We perform an initial classification of each galaxy using its known spectroscopic redshift from a training set (here, the SHELS redshifts) to the best-fitting template from the six default templates for older versions of BPZ, consisting of the four CWW SEDs, plus the Kinney et al. (1996) SB3 and SB2 “starburst” templates (collectively we will refer to this set as the CWW+SB template set), and assign the scaled flux in each filter to the wavelength

$$\lambda_{eff,z} = \frac{\lambda_{eff}}{1 + z_s} \quad (7)$$

Individually, the measurements are quite noisy. We take the mean flux in bins of width $\approx 75-100\text{\AA}$ to obtain a moderate-resolution adjusted SED for each initial template type. If fewer than ten flux points are present in the bin we linearly interpolate between adjacent bins.

Although the procedure could be performed iteratively, reclassifying the training set with the modified SED set, in practice the classifications converge very quickly, and the resultant SEDs are nearly unchanged on subsequent iterations. We do not tweak the two starburst templates, both because we have very few blue galaxies in the magnitude limited SHELS sample, and because our wide bins and broad filters would smooth out emission lines present in the spectra. The SHELS data only cover $0 \leq z \lesssim 0.8$ and given the effective wavelengths of our filter set we can only tweak the wavelength range $\sim 2500-8500\text{\AA}$. Beyond this range we simply use the original template SED, appropriately scaled to match the tweaked portion of the spectrum.

Figure 3 shows the modified SEDs compared to both the original CWW+SB templates and the templates from Coe et al. (2006). One natural consequence of taking median values of the fluxes measured in our broad-band filters, which effectively act as broad smoothing kernels, is

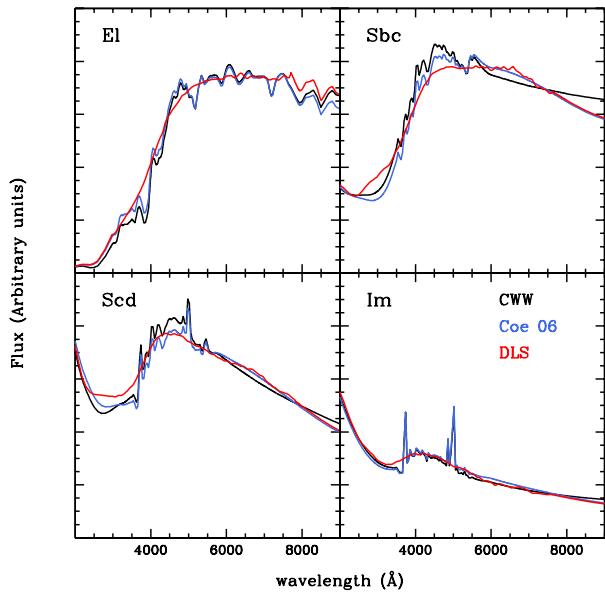


Figure 3. Comparison of the original CWW templates (black) to the DLS tweaked templates (red). The templates of Coe et al. (2006) are shown in blue for comparison. While the modified templates smooth over some small-scale spectral features, the resulting colours are a better match to those observed by DLS for the SHELS data set.

the loss of small scale spectral features of the original templates. While this results in some loss of rigor in our templates, they will have better predicted colours compared to the CWW+SB templates, resulting in less bias and scatter in our photometric redshift predictions. Figure 3 also shows that our templates agree well qualitatively with those of Coe et al. (2006), particularly at the red end.

4.2 Prior Determination

The BPZ software package used to calculate photometric redshifts for DLS includes the application of an apparent-magnitude-based Bayesian prior on both galaxy type and redshift.

The default prior is based on a small set of data from the Hubble Deep Field North (HDFN; Williams et al. 1996). The data set used to train the prior consists of 737 spectroscopic redshifts over the relatively small area of the HDFN, where sample variance may be a problem. It also treats all galaxies brighter than $m < 20$ as if they were 20th magnitude, which can lead to bias for bright objects. In order to examine differences from the HDFN prior and obtain a more accurate prior at brighter magnitudes, we fit a new prior to the same parametrization used by BPZ, using both the DLS SHELS data and the public VIMOS VLT Deep Survey (VVDS) deep spectroscopic survey (Le Fèvre et al. 2005). We use the DLS data set to constrain the prior for galaxies brighter than $R < 21.25$ (Vega) and the VVDS Deep first epoch release catalogue contains 11564 spectra in a 0.61 deg^2 area. While still limited in terms of areal

coverage and subject to sample variance, this data set provides a larger area with more spectroscopic redshifts than the HDFN. After cuts for quality flags and spatial matching we are left with 8145 galaxies, which are assigned a best-fitting template by fixing the correct spectroscopic redshift and finding the best-fitting UBVRI (McCracken et al. 2003) colours.

We use the DLS R -band as the reference magnitude and, following the convention of the DLS, we work in Vega magnitudes, as this is the system that best matches the BVR filter set. In order to place the VVDS r -band magnitudes on the same zero-point scale as the DLS, we cross match VVDS with the r -band magnitudes from CFHTLS Field D1, as CFHTLS r -band has the same zero-point as the SDSS r -band⁴. A simple cross-match of the surveys gives $r_{CFHT} \approx r_{VVDS} + 0.075$ and $r_{SDSS} \approx R_{DLS} + 0.14$ (which includes the 0.20 AB-to-Vega conversion factor for r -band), for an approximate relation $r_{VVDS} \approx R_{DLS} + 0.065$, which we incorporate into the value used for the VVDS m_0 value in the prior.

Benítez (2000) assumes independence of the type and redshift portions of the prior:

$$p(z, T|m_0) = p(T|m_0)p(z|T, m_0) \quad (8)$$

where T is the galaxy type, divided into three broad types as in BPZ as Elliptical (the “El” template), Spiral (the “Sbc” and “Scd” templates), and Irregular (the “Im”, “SB3”, and “SB2” templates), and m_0 is the apparent magnitude of the galaxy in a reference band. The galaxy type fraction is parametrized as an exponential in type fraction as a function of apparent magnitude:

$$p(T|m) = f_T e^{-k_T(m - m_0)} \quad (9)$$

where f_T is the fraction of galaxies of type T at the reference magnitude, m_0 . The type redshift prior is parametrized as:

$$p(z|T, m) \propto z^{\alpha_T} \exp \left\{ - \left(\frac{z}{z_{0T} + k_{mT}(m - m_0)} \right)^{\alpha_T} \right\} \quad (10)$$

where α_T is related to how peaked the redshift distribution of type T is, z_{0T} and k_{mT} control how quickly the peak redshift changes as a function of magnitude, and the proportionality is used to scale the integral of the summed probabilities to unity for proper normalization. To ensure proper normalization, k_T is undefined for the Im/SB type, and is assumed to equal the fraction unassigned to the El and Sp types. For further detail on the prior form, see Benítez (2000).

By parameterizing the prior with two data sets (SHELS for $R < 21.25$ and VVDS for $R \geq 21.25$) we do introduce a possible discontinuity in the prior prediction at $R = 21.25$. We have taken care to minimize this effect by constraining the initial fractions, f_T , of VVDS to match the parametrized fraction predicted from SHELS. Table 1 lists the best-fitting parameters for both the SHELS and VVDS priors. The fit is performed with a basic downhill simplex algorithm, and errors are estimated at the minimum assuming a multidimensional Gaussian in likelihood. The parameters k_T and f_T are fit simultaneously, as are α_T , z_{0T} and k_{mT} . While the best-fitting parameters are a good description of the data, we caution that the likelihood maximization occurs in very

⁴ See http://terapix.iap.fr/rubrique.php?id_rubrique=252

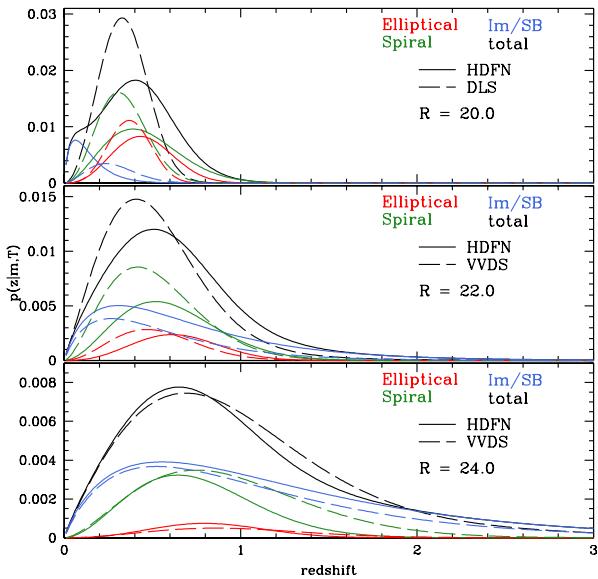


Figure 4. Comparison of the HDFN prior (solid lines) to the DLS/VVDS prior (dashed lines) for Elliptical (red), Spiral (green), Im/SB (blue) and total (black). The priors are markedly different at bright magnitudes, but are similar at $R = 24$. The differences at bright magnitudes may be due to sample variance or the extra freedom afforded by fitting to two data sets.

large five- and nine-dimensional likelihood spaces, so the errors are almost certainly underestimated. Figure 4 shows a comparison of the HDFN prior to the DLS/VVDS prior at three representative magnitudes. The priors are markedly different at bright apparent magnitudes; for instance, the VVDS sample predicts a larger fraction of spiral galaxies compared to HDFN. This may be due to the extra freedom allowed in our prior when we fit $R < 21.25$ with DLS SHELS data and $R > 21.25$ with VVDS data, giving two times the number of free parameters as were used in the HDFN prior. It may also be due to sample variance or Poisson fluctuations in the number of bright galaxies in the limited area of the two spectroscopic surveys in question, particularly the very small HDFN. At $R = 24$ the two priors are quite similar; however, the VVDS prior has a slightly larger tail at high redshift than HDFN, and this difference grows at fainter magnitudes. Once again, this could be due to sample variance or to the extra freedom in the DLS/VVDS parametrization. We will discuss further investigation of the prior in Section 6.

5 PHOTOMETRIC REDSHIFT RESULTS

In order to test our photometric redshift results with the new templates and prior, we require a completely independent test sample with known spectroscopic redshifts. This is normally done by setting aside randomly selected galaxies from the training sample that are not used in the training procedure and calculating their redshift probabilities. While this procedure may be valid, it does result in a test sample that is subject to the same selection criteria, photometric offsets, and other possible systematic effects as the training

sample. A completely independent test set is more desirable, provided that it is complete to sufficient magnitude limits.

The PRIMUS spectroscopic data set described in Section 2.2 provides exactly this, with 9107 spectroscopic redshifts, 703 of which are at $R > 23.3$, measured in a field that is spatially distinct from those used in the training of the templates and prior. As a further benefit, this sample resides in DLS Field F5, observed with the Blanco telescope at CTIO, whereas the SHELS sample used to train both the SEDs and prior is located in Field F2, observed with The Mayall telescope at KPNO. The distinct nature of these two samples enables us to test the photo- z performance in the presence of any field-to-field systematics that might be present in the survey.

Figure 5 shows two contour plots of the probability density $p(z)$ as a function of spectroscopic redshift for the 9107 PRIMUS galaxies in our test sample. Fifty linearly spaced contours are shown, encoding both the density of spectroscopic objects at each redshift and the spread in the PDFs. Red dashed lines indicate the threshold that defines “catastrophic outliers” where the photo- z differs from the spec- z by $\pm 0.15(1+z)$. The left panel uses the unmodified CWW+SB templates and default HDFN prior, while the right panel uses the tweaked templates discussed in Section 4.1 and the DLS/VVDS prior discussed in Section 4.2.

Also shown for comparison are binned averages for two “point” estimates of photo- z : triangles represent “ z best” (Z_B in BPZ), the peak likelihood value for each galaxy, while circles represent “ z average,” the weighted mean value of $p(z)$ for each galaxy. We see excellent agreement between the spectroscopic redshifts and all three photometric redshift estimators for $z > 0.3$. Our shortest-wavelength band, B , with its effective wavelength of 4420Å, does a poor job of constraining the 4000Å break (the dominant feature for photo- z determination) at low redshift. This deficiency manifests itself as an increase in photo- z scatter at $z < 0.25$ for the $p(z)$ distribution. The single-point estimators show a pronounced bias at low redshift, as a single number is unable to capture the the asymmetric PDFs for these galaxies. We strongly discourage use of the single point “ z best” estimators (see, e. g., Abrahamse et al. 2011).

The quality of photometric redshifts is often summarized by three parameters: the (assumed Gaussian) “width” of the distribution (σ_z), the redshift bias, and the catastrophic outlier fraction. We define the redshift bias as the mean value of the difference between the photo- z and spec- z estimates. The fraction of “catastrophic outliers” is defined as the fraction of galaxies with redshift estimates different from the spectroscopic redshift greater than some fixed amount. For the point estimate photo- z ’s, σ_z uncertainty is calculated as:

$$\frac{\sigma_z}{1+z} = \sqrt{\frac{1}{N} \sum_{i=1}^N \frac{(z_{p,i} - z_{s,i})^2}{(1+z_{s,i})^2}} \quad (11)$$

and the catastrophic outlier fraction is defined as the fraction of galaxies with $|z_p - z_s| > 0.15(1+z_s)$. For the full $p(z)$, to measure $\sigma_z/(1+z)$ we sum $p(z) - z_s$ for all galaxies in a redshift interval and fit a Gaussian of width σ . For the catastrophic outlier fraction, we simply sum the amount of probability outside of $|0.15(1+z_s)|$.

Figure 6 shows summary statistics in spectroscopic red-

Table 1. Best-Fitting Prior Parameters

Type	f_T	k_T	α_T	z_{0T}	k_{mT}
DLS $m_0 = 18.0$					
El	0.55 ± 0.02	0.25 ± 0.010	2.99 ± 0.091	0.191 ± 0.0050	0.089 ± 0.0033
Sp	0.39 ± 0.02	-0.175 ± 0.011	2.15 ± 0.058	0.121 ± 0.0055	0.093 ± 0.0040
Im/SB	0.06 ± 0.02		1.77 ± 0.11	0.045 ± 0.0065	0.096 ± 0.014
VVDS $m_0 = 21.25$					
El	0.25 ± 0.02	0.565 ± 0.028	1.957 ± 0.165	0.321 ± 0.028	0.196 ± 0.016
Sp	0.61 ± 0.02	0.155 ± 0.013	1.598 ± 0.08	0.291 ± 0.016	0.167 ± 0.010
Im/SB	0.14 ± 0.02		0.964 ± 0.045	0.170 ± 0.012	0.129 ± 0.013

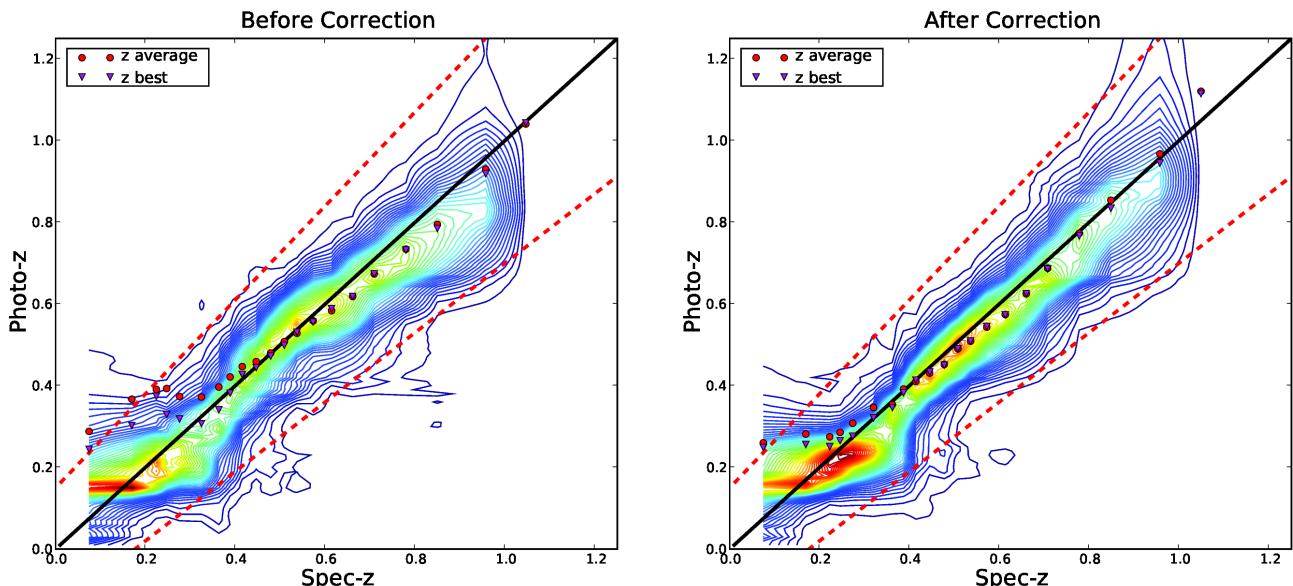


Figure 5. Plot of the summed probability distribution functions of the DLS PRIMUS galaxies as a function of spectroscopic redshift using the original CWW+SB filters and HDFN prior (left) and the tweaked filters and DLS/VVDS prior (right). Contours are linearly spaced and reflect both the number of galaxies and the probability in each redshift interval. Also shown are the mean single point photo-z estimates in twenty bins for both the peak likelihood estimate (z best, purple triangles) and average probability per galaxy (z mean, red circles). Photo-z bias and scatter are reduced by the template and prior optimizations. The threshold that defines the catastrophic outlier fraction are shown as red dashed lines defined by $0.15(1+z)$. The optimization reduces overall redshift scatter by ≈ 20 per cent at $z < 0.6$ and reduces bias by a factor of two at $z > 0.6$. In the improved results the point estimators show a pronounced bias at $z < 0.25$ though this bias can be mitigated by using the full photo-z PDF.

shift bins for the PRIMUS dataset calculated for photo-z's generated with the original CWW+SB templates and HDFN prior and the updated templates and DLS/VVDS prior. The left panels show estimates based on the redshift PDFs, right panels show estimates using the Z_B point estimate. Use of the updated prior and template set led to a ~ 20 per cent reduction in photo-z scatter at $z < 0.6$ and a reduction in photo-z bias by a factor of two at $z > 0.6$. All estimators show a large bias (> 0.06) at $z < 0.2$. This is not unexpected, as we do not allow the photo-z estimators to assign redshifts below $z_p < 0.01$, which naturally leads to a truncation bias at low redshift. The lack of broad spectral features (e.g., the 4000Å break) and the presence of degeneracies in the DLS colours (which will be discussed in the

following subsection) both contribute to this bias, as well as to larger uncertainties in the photo-z estimates. Use of the full $p(z)$ does result in better performance than the point estimate in the $0 \leq z \leq 0.2$ range, with an average uncertainty of 0.064 for $p(z)$ versus 0.102 for the point estimates. However, due to the large scatter we do not recommend the use of DLS photometric redshifts below $Z_B \lesssim 0.25$. The catastrophic outlier rate remains nearly unchanged, as the effects of reduced bias and scatter occur largely within the very broad interval defined by our $> 0.15(1 + z_s)$ catastrophic outlier criterion.

Measuring the resulting scatter in the photo-z versus spec-z relation yields an estimate for the actual uncertainty in measuring photometric redshifts (at least for those ob-

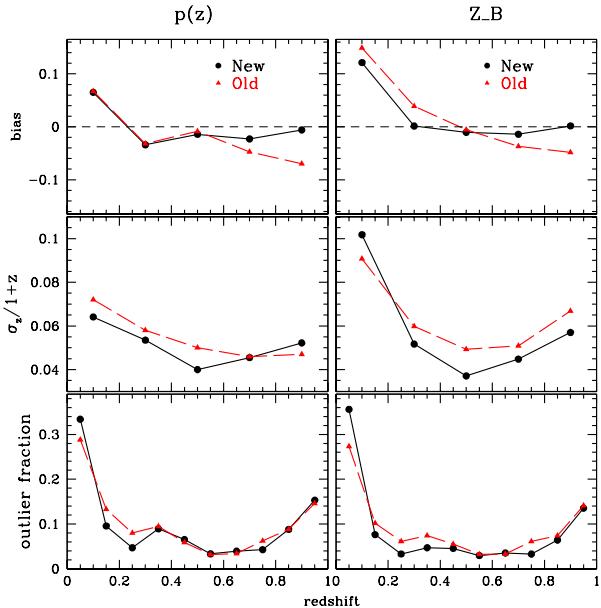


Figure 6. Measured photo-z bias (top panels), photometric scatter ($\sigma_z/(1+z)$, middle panels) and catastrophic outlier fraction (bottom panels) as a function of spectroscopic redshift for the PRIMUS galaxies in DLS computed for both the original CWW+SB templates and HDFN prior (red) and tweaked templates and DLS/VVDS prior (black). The left panels show results using the full $p(z)$ and the right panels show the Z_B point photo-z estimator. Improvements in the templates and prior reduce scatter by 20 per cent and dramatically reduce bias at higher redshift.

jects in the test sample). However, this statistical measure does not tell us anything about the reliability of the redshift probability densities estimated for each galaxy. In calculating the redshift PDF we assume, among other things, that our template set is complete and spanning (and, in nearly all current photo-z codes, that the templates do not evolve). We also assume that our prior is correct, even when extrapolated to fainter magnitudes than are covered by the training set. The breakdown of any of these assumptions can lead to an incorrect estimate of the redshift PDF. In other words, we want to ask: do the redshift uncertainties based on the posterior probability distributions for individual galaxies match the actual measured uncertainties?

To test the accuracy of our $p(z)$ estimates for the PRIMUS data, we determine the redshift bounds containing a fixed fraction of each individual galaxy's total probability (extending from "zmin" to "zmax" and centred on the peak probability, Z_B) and compare this to the actual fraction of spectroscopic redshifts that fall within this interval. This is very similar to a procedure described in Fernández-Soto et al. (2002). Table 2 lists the probability fraction and the number and fraction of galaxies that fall within the corresponding "zmin" to "zmax" interval for all PRIMUS galaxies between $0.01 \leq z \leq 1.0$. If the $p(z)$ estimates are accurate then we expect these fractions to agree. For example, if our 1σ estimates determined from $p(z)$ are correct, then 68 per cent of the spectroscopic redshifts should fall within these 1σ bounds. The values are somewhat noisy, as the DLS photo-z algorithm samples the redshifts at

Table 2. Predicted and Observed Probability Density Fractions

$p(z)$	Fraction	N in interval	Observed Fraction
	0.30	3509/9011	0.39
	0.50	5419/9011	0.60
	0.68	6846/9011	0.76
	0.90	8014/9011	0.89
	0.95	8199/9011	0.91
	0.99	8320/9011	0.92

intervals of $\Delta z = 0.01$, so $p(z)$ is often localized to a small number of non-zero values. To mitigate this resolution effect, we linearly interpolate the PDF to a grid of $\Delta z = 0.0025$. We test the addition of a redshift cut of $0.25 \leq z \leq 1.0$ to avoid the biased $p(z)$ predictions at low redshift. We also select galaxies based on the BPZ ODDS parameter, which is a measure of the amount of probability within a fixed interval of $0.06(1+z)$ around the Z_B value. BPZ ODDS selects galaxies with peaked and unimodal $p(z)$ distributions. The observed fractions when restricting to $0.25 \leq z \leq 1.0$ and $\text{BPZ ODDS} > 0.6$ differ by less than one per cent compared to the total sample shown in Table 2. There is some evidence that we are underestimating the core of the distribution (as 60 per cent of galaxies actually lie within the 50 per cent confidence intervals), and are overconfident in our estimate of the distribution tails (e.g., the 99 per cent confidence region contains only 92 per cent of galaxies). However, Table 2 shows that the $p(z)$ fractions agree to within 10 per cent with the number of spec-z objects and thus the $p(z)$ are approximately representative of the true redshift uncertainties for this data set.

5.1 Type Dependence

We now examine the photo-z's of the PRIMUS sample, broken up by galaxy type. We do not attempt to classify the galaxies from the spectra themselves; rather, we use the best-fitting type found with BPZ's *ONLY_TYPE* option briefly described in Section 4. Figure 7 shows the summed PDFs as a function of PRIMUS spectroscopic sample from $0 \leq z \leq 1$ broken into types. As we run BPZ with *INTERP=2*, each type also includes the adjacent interpolated templates. The starburst templates, SB3 and SB2, appear particularly noisy, partially due to a smaller number of galaxies (only 472 and 668 galaxies classified as SB3 and SB2 respectively) although their less extreme colours, due to a smaller 4000Å break, also lead to increased scatter.

The pronounced bias and larger uncertainties at $z < 0.2$ are noticeable for the Elliptical and Spiral galaxies. These are due to colour degeneracies in the DLS filter set. Figure 8 shows the observed $(B-V)$ (top panels) and $(V-R)$ (bottom panels) colours as a function of spectroscopic redshift (left panels) and photometric redshift (right panels) for the PRIMUS galaxies. The red, green, magenta, and blue lines show the E, Sbc, Scd, and Im model template colours respectively. Green points highlight the galaxies classified as Sbc (left). As expected, the Sbc galaxies match the predicted colours to the best-fitting type at the assigned point photometric redshift very well. However, several prominent areas of departure of the Sbc galaxies from the proper colour track

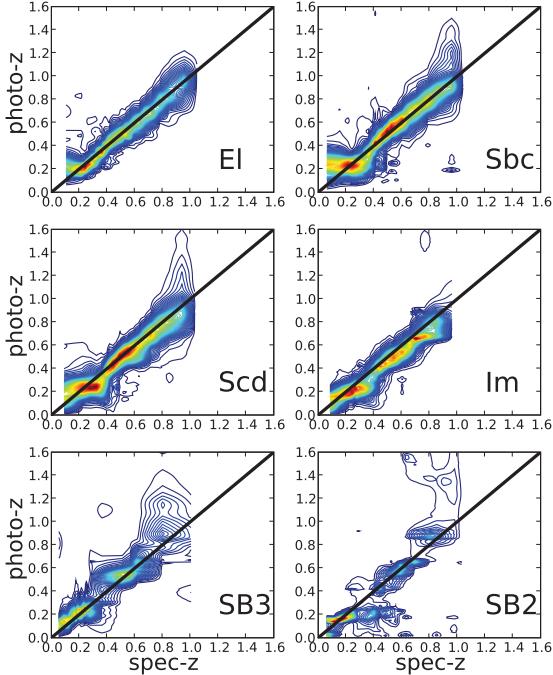


Figure 7. The summed photometric redshift PDFs as a function of spectroscopic redshift for each of the six galaxy templates. The Elliptical and Spiral types show a strong bias at $z < 0.25$ but excellent performance at intermediate redshifts.

are present, particularly the ($B - V$) colour at $z < 0.2$, and almost no galaxies are assigned to types E, Sbc, or Scd at $z < 0.2$. For the DLS filter set, there is a degeneracy between expected colours of the templates at $z \sim 0.0 - 0.1$ and the next-bluer template at $z \sim 0.2 - 0.3$. A particular degeneracy is highlighted by the black and white dots, which show that the Elliptical at $z = 0.06$ and Sbc at $z = 0.23$ have nearly identical colours. ($R - z$) is omitted for brevity, but does not offer much discriminating power, with template values of $(R - z) = 0.59$ for the El and $(R - z) = 0.57$ for the Sbc at the redshifts in question). This is even more evident in Figure 9, which shows the ($V - R$) versus ($B - V$) colour-colour plot for the PRIMUS galaxies in DLS. Model colours for El (red), Sbc (green), Scd (magenta) and Im (blue) are shown from $z = 0.0$ to $z = 1.1$. There is significant overlap in the model colour tracks for the E, Sbc, and Scd galaxies, particularly at low redshift. The degeneracy from Figure 8 is indicated by the black and white dot. Such overlaps in colour-colour space show that the mapping from colour to redshift is not unique, leading to degenerate solutions and broad redshift PDFs. The addition of extra colours (for instance, adding u -band or near-infrared data to the DLS) could break such degeneracies and improve low redshift photo-z results.

5.2 The DLS Galaxy Redshift Distribution

Rather than restricting ourselves to a spectroscopic subsample, we now show samples for the full DLS data set, similar to what would be used in a typical science case, though the ex-

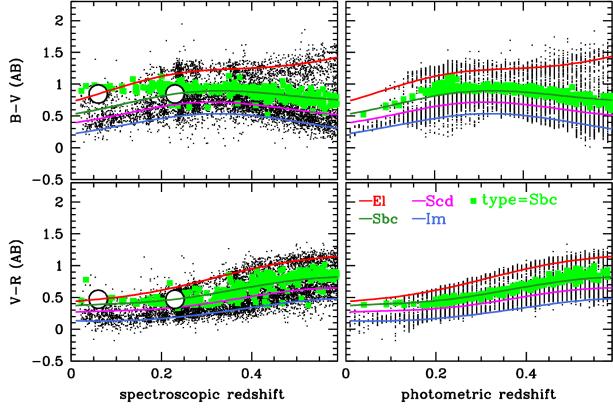


Figure 8. Observed $B - V$ (top) and $V - R$ (bottom) colours as a function of spectroscopic redshift (left) and photometric redshift (right) for the PRIMUS galaxies. Model template colour tracks are shown for the El (red), Sbc (green), Scd (magenta) and Im (blue) templates. The green points show galaxies classified as type Sbc, which should match the green Sbc colour track. Departures from the expected colour tracks are evident in the spectroscopic redshift plot, and almost no E, Sbc, and Scd galaxies are assigned to $z < 0.2$ due to colour degeneracies. A specific degeneracy between El at $z = 0.06$ and Sbc at $z = 0.23$ is highlighted with black and white dots.

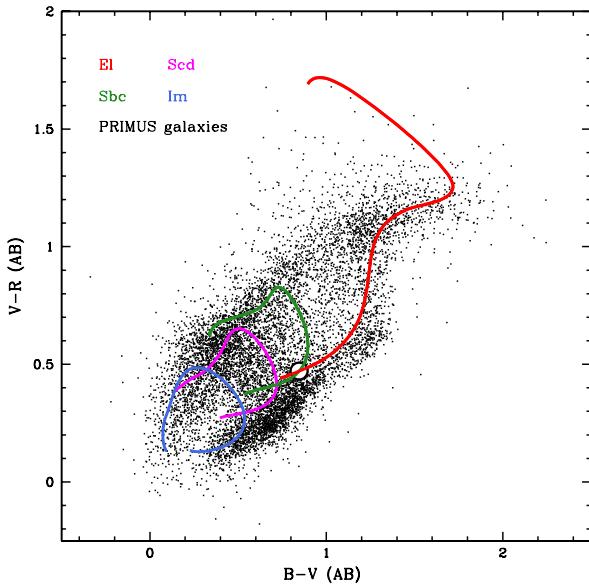


Figure 9. Colour-colour plot for the PRIMUS galaxies in DLS. Model template colour tracks are shown for the El (red), Sbc (green), Scd (magenta) and Im (blue) templates for $z = 0.0$ to $z = 1.1$. The areas of overlap between the model tracks show that significant degeneracies exist at low redshift for the DLS filter set. The particular degeneracy highlighted by the grey stars in Figure 8 is indicated here by the black and white dot at $(B - V) = 0.85$ and $(V - R) = 0.47$.

act selection will depend on the specific needs of the analysis. These photometric redshifts have been used in several DLS-based science results, and each specific data set was subject to tests for systematic errors; see the individual papers for details (Morrison et al. 2012; Jee et al. 2012; Choi et al. 2012). Figure 10 shows the $N(z)$ for apparent-magnitude-limited slices consisting of all five DLS fields. The sample has been trimmed based on bright object and bad pixel masks, and a star/galaxy cut has been made based on the DLS *dlsqcprob* parameter. No cuts have been made on the BPZ ODDS parameter, which measures the amount of posterior probability within a fixed distance of the peak. The solid histogram represents the single point Z_B estimator, while the dashed curve shows the sum of the photo-z PDFs. Comparison of the point estimate to the summed p(z) estimate shows qualitative agreement between the point and p(z) estimators for bright magnitudes, but increasing differences as fainter apparent magnitudes are included, particularly when the sample extends to $z \geq 2$. The excess probability at $z \geq 2$ is due to both broadly peaked PDFs for galaxies at $z \sim 2.0$ and contributions of secondary “catastrophic outlier” peaks in PDFs at $z \sim 0.2 - 0.4$ where the colour degeneracy between the Lyman break and 4000Å break produces bimodal probability distributions. As samples push fainter in apparent magnitude and include more broad-tailed and multi-modal PDFs, use of the full p(z) distribution is essential to obtaining accurate redshift estimates. As is discussed in Jee et al. (2012), using the sum of p(z) values also results in a much smoother form for $N(z)$ that is well-defined on the grid, obviating the need to fit a parametrized form for $N(z)$ that may be necessary if using a point estimate for some science applications.

Figure 10 also shows a distinct bimodal shape for the $N(z)$. The excess probability at $z \sim 0.25$ is due to the degeneracies shown in Figure 8, where galaxies at lower redshift are degenerate with $z \sim 0.25$ galaxies with slightly bluer colours, similar to the degeneracy discussed in Figure 8. This feature is specific to the *BVRz* filter set and lack of constraint on the 4000Å break at low redshift. Even use of p(z) for these galaxies does not fully correct the problem. We once again caution against using DLS photometric redshifts below $Z_B \lesssim 0.25$.

6 CONCLUSIONS AND FUTURE WORK

To optimize the quality of the photo-z’s, we have improved several of the inputs to the photo-z calculation, including the response curves, the prior probabilities, and the spectral templates. It is very important that proper photometry be obtained before adjusting the template SEDs, fitting for a new prior, or calculating photo-z’s. Any bias in the input photometry will mix systematic effects with the underlying relations that we are attempting to recover, often for specific subsets of the data. For example, Figure 1 shows that using the incorrect *z*-band filter curve gives accurate expected colours for blue stars, but results in incorrectly predicted *z*-band magnitudes of more than 0.1 magnitudes for red stars. Any such bias will propagate through the calculation in the form of both misestimated observed fluxes and incorrect galaxy type classifications and can adversely affect photo-z performance. Thus, careful photometric cali-

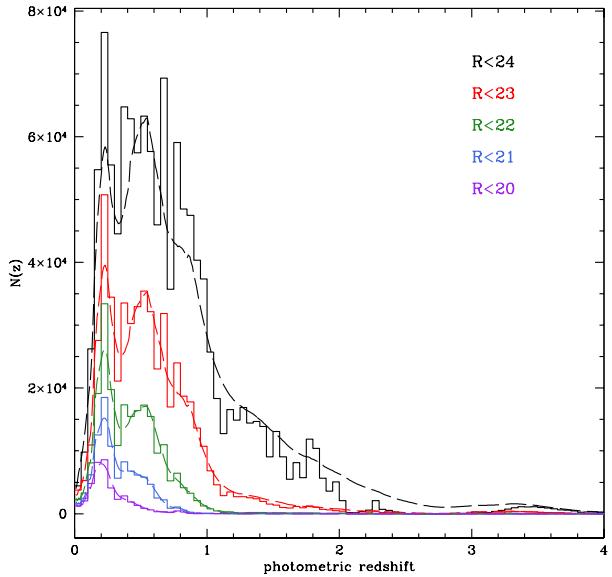


Figure 10. Photometric redshift $N(z)$ summed over all five DLS fields as a function of R magnitude. Solid histogram shows the single point Z_B estimator, while the dashed curves show the sum of the p(z) estimates. The summed p(z) distributions provide a smooth representation of $N(z)$ and correctly account for the long-tailed and bimodal shape of the posterior distributions, particularly at $z > 2$.

bration is of the utmost importance for photometric redshift surveys.

By replacing the published filter+CCD response for the MOSAIC camera’s *z* band with a new CCD response based on the absorption properties of silicon, we improved the predictions of red star colours by 0.1 magnitudes relative to SDSS. Importantly, this was accomplished with a complete revision of the response curve, to be used on all objects in the DLS, rather than a linear colour correction, which would have been calculated using only stars, and may not have fully corrected galaxy photometry. The filter curve is also used in computing the expected model fluxes in the photo-z algorithm, so that by correcting the filter curve itself we have eliminated a troublesome source of bias that was present in the earlier DLS photo-z calculations. We provide the transmission curve electronically for use with other archival data which have used the *z*-band filter with the original MOSAIC cameras at KPNO and CTIO. We tested the improvement in photo-z performance using galaxies from DLS field F5, which overlaps the PRIMUS survey, and found a 20 per cent improvement in photo-z scatter at $z < 0.6$ and a reduction in photo-z bias of more than a factor of two at $z > 0.6$ (Figure 6). Catastrophic outlier rates remained nearly unchanged.

One way of further reducing photo-z errors is by adding flux measurements in different passbands, beyond the original four DLS bands. The addition of shorter wavelength bands (*u*-band, for example) that bracket the 4000Å break at low redshift can break degeneracies of the type described in Section 5.1. The near-infrared is also helpful in disambiguating galaxy types, so the addition of, e.g., *JHK_S* data

should also improve photo-z performance at low redshift, and the detection of the 4000Å break in the near-infrared can reduce the rate of catastrophic photo-z errors that result from misidentifying the Lyman break. We have observed $\simeq 8$ deg 2 of the DLS area in the J and K_s bands using the NEWFIRM detector on the CTIO and KPNO 4-m telescopes and the WIRC instrument on the Palomar 5-m. We will examine the effect of including these bands in a future paper. One DLS field (F2) has been observed in the mid-IR using IRAC 4-band photometry, which will also be examined in the future; however, additional work to improve the mid-IR templates will be required in order for the photo-z's to be improved rather than degraded by the addition of IRAC imaging (see, e.g. Hildebrandt et al. 2010).

Lacking spectroscopic redshifts in our training and test sets at $z > 1$ we cannot evaluate photo-z performance in the DLS at high redshift. However, there are several methods that can illuminate expected photo-z results. We can compare our photo-z $N(z)$ to deeper surveys such as COSMOS that have both deeper spectroscopic samples and a larger number of broad-band observations which produce higher precision photo-z's. Using only bands similar to our $BVRz$ filter complement will enable us to estimate expected bias, scatter, and catastrophic outlier rates at fainter magnitudes. Beyond direct comparison to representative spectroscopic samples, we can use spatial cross-correlations with a bright subset of galaxies (e.g., Matthews & Newman 2010, Schmidt et al. submitted) to measure the redshift distribution of the survey. However, the spectroscopic samples in question must span the redshift range of interest. With both the SHELS and PRIMUS samples restricted to $z \lesssim 1$ we require further targeted spectroscopy in the DLS to fully utilize this technique.

As discussed earlier in the paper, much photometric redshift work to date has assumed that the template SEDs do not evolve and has used discretized types. A few examples exist for fitting a continuous galaxy type (Blanton & Roweis 2007; Brammer et al. 2008), usually with linear combinations of a set of basis templates. This form naturally allows for changes in the average template with redshift, but makes type dependent Bayesian priors more difficult to implement. Methods of parametrizing template evolution and continuous galaxy type while simultaneously allowing for type-dependent priors will be explored in the future (Schmidt et al. in prep.). However, a larger number of narrower filter curves would be preferred for detailed examination of galaxy types and evolution. Surveys such as J-PAS (Benítez et al. 2009) would be ideal for the study of template and prior evolution.

ACKNOWLEDGEMENTS

The authors would like to thank David Wittman, Tony Tyson, Dinesh Loomba, James Jee, Russell Ryan, Chris Morrison, Ami Choi, Vera Margoniner, Begoña Ascaso, and Jim Bosch for useful discussions and suggestions during the work discussed in this paper.

SJS was supported by NSF Grant AST-1009514.

PAT was partly supported by a grant from the New Mexico Space Grant Consortium.

Funding for the Deep Lens Survey has been provided

by Bell Labs Lucent Technologies and NSF grants AST 04-41072 and AST 01-34753. Observations were obtained at Cerro Tololo Inter-American Observatory and Kitt Peak National Observatory. CTIO and KPNO are divisions of the National Optical Astronomy Observatory (NOAO), which is operated by the Association of Universities for Research in Astronomy, Inc., under cooperative agreement with the National Science Foundation.

We would like to thank the SHELS team for making a portion of their spectroscopic sample available to us for use in the photometric redshift analysis.

We thank the PRIMUS team for sharing their redshift catalogue. Funding for PRIMUS has been provided by NSF grants AST-0607701, 0908246, 0908442, 0908354, and NASA grant 08-ADP08-0019. This paper includes data gathered with the 6.5-metre Magellan Telescopes located at Las Campanas Observatory, Chile.

We thank CFHTLS for making their data available, observations obtained with MegaPrime/MegaCam, a joint project of CFHT and CEA/DAPNIA, at the Canada-France-Hawaii Telescope (CFHT) which is operated by the National Research Council (NRC) of Canada, the Institut National des Sciences de l'Univers of the Centre National de la Recherche Scientifique (CNRS) of France, and the University of Hawaii. This work is based in part on data products produced at TERAPIX and the Canadian Astronomy Data Centre as part of the Canada-France-Hawaii Telescope Legacy Survey, a collaborative project of NRC and CNRS.

REFERENCES

- Abrahamse A., Knox L., Schmidt S., Thorman P., Tyson J. A., Zhan H., 2011, ApJ, 734, 36
- Aihara H. et al., 2011, ApJS, 193, 29
- Benítez N., 2000, ApJ, 536, 571
- Benítez N. et al., 2009, ApJ, 691, 241
- Blanton M. R., Roweis S., 2007, AJ, 133, 734
- Bolzonella M., Miralles J.-M., Pelló R., 2000, A & Ap., 363, 476
- Brammer G. B., van Dokkum P. G., Coppi P., 2008, ApJ, 686, 1503
- Bruzual, Persson, Gunn, Stryker, retrieved 10/06/2009, StSCI Calibration Database
- Budavári T., Szalay A. S., Connolly A. J., Csabai I., Dickinson M., 2000, AJ, 120, 1588
- Choi A., Tyson J. A., Morrison C. B., Jee M. J., Schmidt S. J., Margoniner V. E., Wittman D. M., 2012, ApJ, 759, 101
- Clampin M. et al., 1998, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 3356, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, P. Y. Bely & J. B. Breckinridge, ed., pp. 332–337
- Coe D., Benítez N., Sánchez S. F., Jee M., Bouwens R., Ford H., 2006, AJ, 132, 926
- Coil A. L. et al., 2011, ApJ, 741, 8
- Coleman G. D., Wu C.-C., Weedman D. W., 1980, ApJS, 43, 393
- Covey K. R. et al., 2007, AJ, 134, 2398
- Fernández-Soto A., Lanzetta K. M., Chen H.-W., Levine B., Yahata N., 2002, MNRAS, 330, 889

- Geller M. J., Dell'Antonio I. P., Kurtz M. J., Ramella M., Fabricant D. G., Caldwell N., Tyson J. A., Wittman D., 2005, ApJL, 635, L125
- Geller M. J., Kurtz M. J., Dell'Antonio I. P., Ramella M., Fabricant D. G., 2010, ApJ, 709, 832
- Hildebrandt H. et al., 2010, A & Ap., 523, A31
- Hildebrandt H. et al., 2012, MNRAS, 421, 2355
- Hinkle K. H., Wallace L., Livingston W., 2003, in Bulletin of the American Astronomical Society, Vol. 35, American Astronomical Society Meeting Abstracts, p. 1260
- Ilbert O. et al., 2006, A & Ap., 457, 841
- Jee M. J., Tyson J. A., Schneider M. D., Wittman D., Schmidt S., Hilbert S., 2012, ArXiv e-prints
- Kaiser N. et al., 2002, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 4836, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, J. A. Tyson & S. Wolff, ed., pp. 154–164
- Kinney A. L., Calzetti D., Bohlin R. C., McQuade K., Storchi-Bergmann T., Schmitt H. R., 1996, ApJ, 467, 38
- Laureijs R. J., Duvet L., Escudero Sanz I., Gondoin P., Lumb D. H., Oosterbroek T., Saavedra Criado G., 2010, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7731, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series
- Le Fèvre O. et al., 2005, A & Ap., 439, 845
- Margoniner V. E., Wittman D. M., 2008, ApJ, 679, 31
- Matthews D. J., Newman J. A., 2010, ApJ, 721, 456
- McCracken H. J. et al., 2003, A & Ap., 410, 17
- Morrison C. B., Scranton R., Ménard B., Schmidt S. J., Tyson J. A., Ryan R., Choi A., Wittman D. M., 2012, MNRAS, 426, 2489
- Padmanabhan N. et al., 2008, ApJ, 674, 1217
- Pickles A. J., 1997, VizieR Online Data Catalog, 7102, 0
- Rajkaran K., Singh R., Shewchun J., 1979, Solid State Electronics, 22, 793
- Ryan R. E. et al., 2011, ApJ, 739, 83
- Scientific Imaging Technologies Inc., 1995, Literature No. ST-002A, 12th edn. Retrieved from <http://www.ociw.edu/instrumentation/ccd/parts/ST-002a.pdf>
- Tyson J. A., 2002, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 4836, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, J. A. Tyson & S. Wolff, ed., pp. 10–20
- Varshni Y. P., 1967, Physica, 34, 149
- Wester W., Dark Energy Survey Collaboration, 2005, in Astronomical Society of the Pacific Conference Series, Vol. 339, Observing Dark Energy, Wolff S. C., Lauer T. R., eds., p. 152
- Williams R. E. et al., 1996, AJ, 112, 1335
- Wittman D., Riechers P., Margoniner V. E., 2007, ApJL, 671, L109
- Wittman D., Ryan R., Thorman P., 2012, MNRAS, 421, 2251
- Wittman D. M. et al., 2002, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, J. A. Tyson & S. Wolff, ed., Vol. 4836, pp. 73–82